



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/668,467	09/23/2003	Sachin Mullick	10830.0100.NPUS00	2942
27927 7590 06/24/2009 RICHARD AUCHTERLONIE NOVAK DRUCE & QUIGG, LLP 1000 LOUISIANA 53RD FLOOR HOUSTON, TX 77002				
EXAMINER				
BIBBEE, JARED M				
ART UNIT		PAPER NUMBER		
2161				
MAIL DATE		DELIVERY MODE		
06/24/2009		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary

Application No.

10/668,467

Applicant(s)

MULLICK ET AL.

Examiner

JARED M. BIBBEE

Art Unit

2161

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 15 April 2009.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-28, 32-54 and 58-73 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-28, 32-54 and 58-73 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☒ Information Disclosure Statement(s) (PTO/SB-08)
Paper No(s)/Mail Date 4/15/2009
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date _____
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: _____

DETAILED ACTION

Response to Amendment

1. This Office Action has been issued in response to appeal brief filed on 15 April 2009. Claims 29-31 and 55-57 are cancelled. Claims 1-28, 32-54, and 58-73 are pending. Applicants' arguments have been carefully and respectfully considered in light of the affidavits and are persuasive, as they relate to the claim rejections under 35 U.S.C. 102 and 103. Therefore, those rejections are withdrawn. However, after further search and consideration, Examiner has found additional prior art, as will be discussed below. Accordingly, this action has been made NON-FINAL.

Claim Rejections - 35 USC § 102

2. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

3. **Claims 1-4, 10, 11, 15, 22, 32, 33-36, 44, 45, 49, 58, and 61-73 are rejected under 35 U.S.C. 102(e) as being anticipated by Burns et al (US 6,925,515 B2).**

With respect to independent claim 1, Burns teaches a method of operating a network file server computer for providing clients with concurrent write access to a file, the method comprising the network file server computer responding to a concurrent write request from a client by:

- a) obtaining a lock for the file (*see column 8, lines 49-67*); and then
- b) preallocating a metadata block for the file (*see column 10, lines 14-30*); and then
- c) releasing the lock for the file (*see column 10, lines 14-30*); and then
- d) asynchronously writing to the file (*see column 10, lines 14-30*); and then
- e) obtaining the lock for the file (*see column 10, lines 31-44*); and then
- f) committing the metadata block to the file in the data storage (*see column 10, line 65 through column 11, line 5*); and then
- g) releasing the lock for the file (*see column 11, lines 6-14*).

With respect to dependent claim 2, Burns further teaches a hierarchy of blocks including an inode block of metadata, data blocks of file data, and indirect blocks of metadata, and wherein the metadata block for the file is an indirect block of metadata (*see column 10, line 14 through column 11, line 14*).

With respect to dependent claim 3, Burns further teaches copying data from an original indirect block of the file to the metadata block for the file, the original indirect block of the file having been shared between the file and a read-only version of the file (*see column 10, line 14 through column 11, line 14*).

With respect to dependent claim 4, Burns further teaches concurrent writing for more than one client to the metadata block for the file (*see column 10, line 14 through column 11, line 14*).

With respect to dependent claim 10, Burns further teaches writing the metadata block to a log in storage of the network file server computer for committing the metadata block for the file (*see column 10, line 14 through column 11, line 14*).

With respect to dependent claim 11, Burns further teaches gathering together preallocated metadata blocks for a plurality of client write requests to the file, and committing together the preallocated metadata blocks for the plurality of client write requests to the file by obtaining the lock for the file, committing the gathered preallocated metadata blocks for the plurality of client write requests to the file, and then releasing the lock for the file (*see column 10, line 14 through column 11, line 14*).

With respect to independent claim 15, Burns teaches a method of operating a network file server computer for providing clients with concurrent write access to a file, the method comprising the network file server computer responding to a concurrent write request from a client by:

- a) preallocating a metadata block for the file (*see column 10, lines 14-30*); and then
- b) asynchronously writing to the file (*see column 10, lines 14-30*); and then
- c) committing the metadata block to the file in the data storage (*see column 10, line 65 through column 11, line 5*);

wherein the method includes gathering together preallocated metadata blocks for a plurality of client write requests to the file, and committing together the preallocated metadata blocks for the plurality of client write requests to the file by obtaining a lock for the file, committing the gathered preallocated metadata blocks for the plurality of client write requests to the file, and then releasing the lock for the file (*see column 10, line 14 through column 11, line 14*).

With respect to dependent claim 22, Burns further teaches processing multiple concurrent read and write requests by pipelining the requests through a first processor and a second

processor, the first processor performing metadata management for the multiple concurrent read and write requests, and the second processor performing asynchronous reads and writes for the multiple concurrent read and write requests (*see column 10, line 14 through column 11, line 14*).

With respect to independent claim 32, Burns teaches A method of operating a network file server computer for providing clients with concurrent write access to a file, the method comprising the network file server computer responding to a concurrent write request from a client by executing a write thread, execution of the write thread including:

- a) obtaining an allocation mutex for the file (*see column 10, lines 14-30*); and then
- b) preallocating new metadata blocks that need to be allocated for writing to the file (*see column 10, lines 14-30*); and then
- c) releasing the allocation mutex for the file (*see column 10, lines 14-30*); and then
- d) issuing asynchronous write requests for writing to the file (*see column 10, lines 14-30*);
- e) waiting for callbacks indicating completion of the asynchronous write requests (*see column 10, lines 14-30*); and then
- f) obtaining the allocation mutex for the file (*see column 10, lines 31-44*); and then
- g) committing the preallocated metadata blocks (*see column 10, line 65 through column 11, line 5*); and then
- h) releasing the allocation mutex for the file (*see column 11, lines 6-14*).

With respect to claims 33-36, claims 33-36 correspond to claims 1-4 and are rejected for the same reasons as set forth in the rejection of claims 1-4 above.

With respect to claims 44-45, claims 44-45 correspond to claims 10-11 and are rejected for the same reasons as set forth in the rejection of claims 10-11 above.

With respect to claim 49, claim 49 corresponds to claim 15 and is rejected for the same reasons as set forth in the rejection of claim 15 above.

With respect to claim 58, claim 58 corresponds to claim 32 and is rejected for the same reasons as set forth in the rejection of claim 32 above.

With respect to independent claim 61, Burns teaches a method of operating a network file server for providing clients with concurrent write access to a file, the method comprising the network file server responding to a concurrent write request from a client by:

- a) preallocating a block for the file (*see column 10, lines 14-30*); and then
- b) asynchronously writing to the file (*see column 10, lines 14-30*); and then
- c) committing the preallocated block (*see column 10, line 65 through column 11, line 5*);

wherein the network file server also includes an uncached write interface, a file system cache, and a cached read-write interface, wherein the uncached write interface bypasses the file system cache for sector-aligned write operations, and the network file server is programmed to invalidate cache blocks in the file system cache including sectors being written to by the uncached write interface (*see column 10, line 14 through column 11, line 14*).

With respect to dependent claim 62, Burns further teaches a final step of returning to said client an acknowledgement of the writing to the file (*see column 10, line 14 through column 11, line 14*).

With respect to claims 63-67, claims 63-67 correspond to claim 62 and are rejected for the same reasons as set forth in the rejection of claim 62 above.

With respect to dependent claim 68, Burns further teaches a final step of saving the file in disk storage of the network file server (*see column 10, line 14 through column 11, line 14*).

With respect to claims 69-73, claims 69-73 correspond to claim 68 and are rejected for the same reasons as set forth in the rejection of claim 68 above.

Claim Rejections - 35 USC § 103

4. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. Claims 5-9, 12-14, 16-21, 23-28, 37-43, 46-48, and 50-54 are rejected under 35 U.S.C. 103(a) as being unpatentable over Burns in view of Marcotte (US 6,449, 614 B1).

With respect to dependent claim 5, note the discussion of claim 1 above, Burns discloses all of the elements of claim 1, but fails to explicitly recite the asynchronous writing to the file includes a partial write to a new block that has been copied at least in part from an original block of the file, and wherein the method further includes checking a partial block conflict queue for a conflict with a concurrent write to the new block, and upon failing to find an indication of a conflict with a concurrent write to the new block, preallocating the new block, copying at least a

portion of the original block of the file to the new block, and performing the partial write to the new block.

However, Marcotte teaches the asynchronous writing to the file includes a partial write to a new block that has been copied at least in part from an original block of the file, and wherein the method further includes checking a partial block conflict queue for a conflict with a concurrent write to the new block, and upon failing to find an indication of a conflict with a concurrent write to the new block, preallocating the new block, copying at least a portion of the original block of the file to the new block, and performing the partial write to the new block (*see column 13, line 35 through column 14, line 43*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the holding queue of Marcotte for the purpose of achieving better performance eliminate the need to suspend the system.

The suggestion/motivation for doing so would have been to manage locks and improve the system (*see column 20, lines 35-49*).

Therefore, it would have been obvious to combine Marcotte with Burns to obtain the invention as specified in the instant claims.

With respect to dependent claim 6, note the discussion of claim 1 above, Burns discloses all of the elements of claim 1, but fails to explicitly recite the asynchronous writing to the file includes a partial write to a new block that has been copied at least in part from an original block of the file, and wherein the method further includes checking a partial block conflict queue for a conflict with a concurrent write to the new block, and upon finding an indication of a conflict

with a concurrent write to the new block, waiting until resolution of the conflict with the concurrent write to the new block, and then performing the partial write to the new block.

However, Marcotte teaches the asynchronous writing to the file includes a partial write to a new block that has been copied at least in part from an original block of the file, and wherein the method further includes checking a partial block conflict queue for a conflict with a concurrent write to the new block, and upon finding an indication of a conflict with a concurrent write to the new block, waiting until resolution of the conflict with the concurrent write to the new block, and then performing the partial write to the new block (*see column 13, line 35 through column 14, line 43*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the holding queue of Marcotte for the purpose of achieving better performance eliminate the need to suspend the system.

The suggestion/motivation for doing so would have been to manage locks and improve the system (*see column 20, lines 35-49*).

Therefore, it would have been obvious to combine Marcotte with Burns to obtain the invention as specified in the instant claims.

With respect to dependent claim 7, Marcotte further teaches placing a request for the partial write in a partial write wait queue upon finding an indication of a conflict with a concurrent write to the new block, and performing the partial write upon servicing the partial write wait queue (*see column 13, line 35 through column 14, line 43*).

With respect to dependent claim 8, note the discussion of claim 1 above, Burns discloses all of the elements of claim 1, but fails to explicitly recite checking an input-output list for a conflicting prior concurrent access to the file, and upon finding a conflicting prior concurrent access to the file, suspending the asynchronous writing to the file until the conflicting prior concurrent access to the file is no longer conflicting.

However, Marcotte teaches checking an input-output list for a conflicting prior concurrent access to the file, and upon finding a conflicting prior concurrent access to the file, suspending the asynchronous writing to the file until the conflicting prior concurrent access to the file is no longer conflicting (*see column 12, line 7 through column 13, line 32*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the waiting list of Marcotte.

The suggestion/motivation for doing so would have been to manage locks and improve the system (*see column 20, lines 35-49*).

Therefore, it would have been obvious to combine Marcotte with Burns to obtain the invention as specified in the instant claims.

With respect to dependent claim 9, Marcotte further teaches providing a sector-level granularity of byte range locking for concurrent write access to the file by the suspending of the asynchronous writing to the file until the conflicting prior concurrent access is no longer conflicting (*see column 12, line 7 through column 13, line 32*).

With respect to dependent claim 12, note the discussion of claim 1 above, Burns discloses all of the elements of claim 1, Burns further teaches checking whether a previous commit is in

progress after asynchronously writing to the file and before obtaining the lock for the file for committing the metadata block to the file (*see column 10, line 14 through column 11, line 14*), but fails to explicitly recite upon finding that a previous commit is in progress, placing a request for committing the metadata block to the file on a staging queue for the file.

However, Marcotte teaches upon finding that a previous commit is in progress, placing a request for committing the metadata block to the file on a staging queue for the file (*see column 12, line 7 through column 13, line 32*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the waiting list of Marcotte.

The suggestion/motivation for doing so would have been to manage locks and improve the system (*see column 20, lines 35-49*).

Therefore, it would have been obvious to combine Marcotte with Burns to obtain the invention as specified in the instant claims.

With respect to independent claim 13, Burns teaches a method of operating a network file server computer for providing clients with concurrent write access to a file in data storage, the method comprising the network file server computer responding to a concurrent write request from a client by:

- a) preallocating a block for the file (*see column 10, lines 14-30*); and then
- b) asynchronously writing to the file (*see column 10, lines 14-30*); and then
- c) committing the block to the file in the data storage (*see column 10, line 65 through column 11, line 5*);

Burns fails to explicitly recite the asynchronous writing to the file includes a partial write to a new block that has been copied at least in part from an original block of the file, and wherein the method further includes checking a partial block conflict queue for a conflict with a concurrent write to the new block, and upon finding an indication of a conflict with a concurrent write to the new block, waiting until resolution of the conflict with the concurrent write to the new block, and then performing the partial write to the new block.

However, Marcotte teaches the asynchronous writing to the file includes a partial write to a new block that has been copied at least in part from an original block of the file, and wherein the method further includes checking a partial block conflict queue for a conflict with a concurrent write to the new block, and upon finding an indication of a conflict with a concurrent write to the new block, waiting until resolution of the conflict with the concurrent write to the new block, and then performing the partial write to the new block (*see column 13, line 35 through column 14, line 43*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the holding queue of Marcotte for the purpose of achieving better performance eliminate the need to suspend the system.

The suggestion/motivation for doing so would have been to manage locks and improve the system (*see column 20, lines 35-49*).

Therefore, it would have been obvious to combine Marcotte with Burns to obtain the invention as specified in the instant claims.

With respect to dependent claim 14, Marcotte further teaches placing a request for the partial write in a partial write wait queue upon finding an indication of a conflict with a concurrent write to the new block, and performing the partial write upon servicing the partial write wait queue (*see column 13, line 35 through column 14, line 43*).

With respect to dependent claim 16, note the discussion of claim 15 above, Burns discloses all of the elements of claim 15, Burns further teaches checking whether a previous commit is in progress after asynchronously writing to the file and before obtaining the lock for the file for committing the metadata block to the file (*see column 10, line 14 through column 11, line 14*), but fails to explicitly recite upon finding that a previous commit is in progress, placing a request for committing the metadata block to the file on a staging queue for the file.

However, Marcotte teaches upon finding that a previous commit is in progress, placing a request for committing the metadata block to the file on a staging queue for the file (*see column 12, line 7 through column 13, line 32*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the waiting list of Marcotte.

The suggestion/motivation for doing so would have been to manage locks and improve the system (*see column 20, lines 35-49*).

With respect to dependent claim 17, note the discussion of claim 15 above, Burns discloses all of the elements of claim 15, but fails to explicitly recite the network file server computer responding to concurrent write requests by writing new data for specified blocks of the

file to the disk storage without writing the new data for the specified blocks of the file to the file system cache, and invalidating the specified blocks of the file in the file system cache.

However, Marcotte teaches the network file server responding to concurrent write requests by writing new data for specified blocks of the file to the disk storage without writing the new data for the specified blocks of the file to the file system cache, and invalidating the specified blocks of the file in the file system cache (*see column 12, line 7 through column 13, line 32*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the waiting list of Marcotte.

The suggestion/motivation for doing so would have been to manage locks and improve the system (*see column 20, lines 35-49*).

With respect to dependent claim 18, Marcotte further teaches the network file server computer responding to read requests for file blocks not found in the file system cache by reading the file blocks from the file system in disk storage and then checking whether the file blocks have become stale due to concurrent writes to the file blocks, and writing to the file system cache a file block that has not become stale, and not writing to the file system cache a file block that has become stale (*see column 12, line 7 through column 13, line 32*).

With respect to dependent claim 19, Marcotte further teaches the network file server computer checking a read-in-progress flag for a file block upon finding that the file block is not in the file system cache, and upon finding that the read-in-progress flag indicates that a prior read

of the file block is in progress from the file system in the disk storage, waiting for completion of the prior read of the file block from the file system in the disk storage, and then again checking whether the file block is in the file system cache (*see column 12, line 7 through column 13, line 32*).

With respect to dependent claim 20, Marcotte further teaches the network file server computer setting a read-in-progress flag for a file block upon finding that the file block is not in the file system cache and then beginning to read the file block from the file system in disk storage, clearing the read-in-progress flag upon writing to the file block on disk, and inspecting the read-in-progress flag to determine whether the file block has become stale due a concurrent write to the file block (*see column 12, line 7 through column 13, line 32*).

With respect to dependent claim 21, Marcotte further teaches the network file server computer maintaining a generation count for each read of a file block from the file system in the disk storage in response to a read request for a file block that is not in the file system cache, and checking whether a file block having been read from the file system in the disk storage has become stale by checking whether the generation count for the file block having been read from the file system is the same as the generation count for the last read request for the same file block (*see column 12, line 7 through column 13, line 32*).

With respect to dependent claim 23, note the discussion of claim 15 above, Burns discloses all of the elements of claim 15, but fails to explicitly recite serializing the reads by delaying access for each read to a block that is being written to by a prior, in-progress write until completion of the write to the block that is being written to by the prior, in-progress write.

However, Marcotte teaches serializing the reads by delaying access for each read to a block that is being written to by a prior, in-progress write until completion of the write to the block that is being written to by the prior, in-progress write (*see column 12, line 7 through column 13, line 32*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the waiting list of Marcotte.

The suggestion/motivation for doing so would have been to manage locks and improve the system (*see column 20, lines 35-49*).

With respect to dependent claim 24, note the discussion of claim 15 above, Burns discloses all of the elements of claim 15, but fails to explicitly recite serializing the writes by delaying access for each write to a block that is being accessed by a prior, in-progress read or write until completion of the read or write to the block that is being accessed by the prior, in-progress read or write.

However, Marcotte teaches serializing the reads by delaying access for each read to a block that is being written to by a prior, in-progress write until completion of the write to the block that is being written to by the prior, in-progress write (*see column 12, line 7 through column 13, line 32*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the waiting list of Marcotte.

The suggestion/motivation for doing so would have been to manage locks and improve the system (*see column 20, lines 35-49*).

With respect to independent claim 25, Burns teaches a method of operating a network file server computer for providing clients with concurrent write access to a file in data storage, the method comprising the network file server computer responding to a concurrent write request from a client by:

- a) preallocating a metadata block for the file (*see column 10, lines 14-30*); and then
- b) asynchronously writing to the file (*see column 10, lines 14-30*); and then
- c) committing the metadata block to the file in the data storage (*see column 10, line 65 through column 11, line 5*);

Burns fails to explicitly recite the network file server computer responding to concurrent write requests by writing new data for specified blocks of the file to the disk storage without writing the new data for the specified blocks of the file to the file system cache, and invalidating the specified blocks of the file in the file system cache.

Burns also fails to explicitly recite the network file server computer responding to read requests for file blocks not found in the file system cache by reading the file blocks from the file system in disk storage and then checking whether the file blocks have become stale due to concurrent writes to the file blocks, and writing to the file system cache a file block that has not become stale, and not writing to the file system cache a file block that has become stale.

However, Marcotte teaches the network file server computer responding to concurrent write requests by writing new data for specified blocks of the file to the disk storage without

writing the new data for the specified blocks of the file to the file system cache, and invalidating the specified blocks of the file in the file system cache (*see column 12, line 7 through column 13, line 32*).

Marcotte also teaches the network file server computer responding to read requests for file blocks not found in the file system cache by reading the file blocks from the file system in disk storage and then checking whether the file blocks have become stale due to concurrent writes to the file blocks, and writing to the file system cache a file block that has not become stale, and not writing to the file system cache a file block that has become stale (*see column 12, line 7 through column 13, line 32*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the waiting list of Marcotte.

The suggestion/motivation for doing so would have been to manage locks and improve the system (*see column 20, lines 35-49*).

With respect to dependent claim 26, Marcotte further teaches the network file server computer checking a read-in-progress flag for a file block upon finding that the file block is not in the file system cache, and upon finding that the read-in-progress flag indicates that a prior read of the file block is in progress from the file system in the disk storage, waiting for completion of the prior read of the file block from the file system in the disk storage, and then again checking whether the file block is in the file system cache (*see column 12, line 7 through column 13, line 32*).

With respect to dependent claim 27, Marcotte further teaches the network file server computer setting a read-in-progress flag for a file block upon finding that the file block is not in the file system cache and then beginning to read the file block from the file system in disk storage, clearing the read-in-progress flag upon writing to the file block on disk, and inspecting the read-in-progress flag to determine whether the file block has become stale due a concurrent write to the file block (*see column 12, line 7 through column 13, line 32*).

With respect to dependent claim 28, Marcotte further teaches the network file server computer maintaining a generation count for each read of a file block from the file system in the disk storage in response to a read request for a file block that is not in the file system cache, and checking whether a file block having been read from the file system in the disk storage has become stale by checking whether the generation count for the file block having been read from the file system is the same as the generation count for the last read request for the same file block (*see column 12, line 7 through column 13, line 32*).

With respect to claims 37-43, claims 37-43 correspond to claims 5-9 and are rejected for the same reasons as set forth in the rejection of claims 5-9 above.

With respect to claims 46-48, claims 46-48 correspond to claims 12-14 and are rejected for the same reasons as set forth in the rejection of claims 12-14 above.

With respect to claim 50, claim 50 corresponds to claim 16 and is rejected for the same reasons as set forth in the rejection of claim 16 above.

With respect to claims 51-54, claims 51-54 correspond to claims 25-28 and are rejected for the same reasons as set forth in the rejection of claims 25-28 above.

6. Claims 59 and 60 are rejected under 35 U.S.C. 103(a) as being unpatentable over Burns in view of Xu et al (US 6,324,581 B1).

With respect to dependent claim 59, note the discussion of claim 58 above, Burns discloses all of the elements of claim 58 but fails to explicitly recite an uncached write interface, a file system cache and a cached read-write interface, and wherein the uncached write interface bypasses the file system cache for sector-aligned write operations.

However, Xu teaches an uncached write interface, a file system cache and a cached read-write interface, and wherein the uncached write interface bypasses the file system cache for sector-aligned write operations (*see column 8, line 36 through column 9, line 39*).

At the time of the invention, it would have been obvious to one of ordinary skill having the teachings of Burns and Marcotte before him or her, to modify the asynchronous writing of Burns to include the uncached write interface of Xu.

The suggestion/motivation for doing so would have been to reduce the loading of data from data movers by using a cached array (*see column 2, lines 57-61*).

With respect to dependent claim 60, Xu further teaches the network file server is further programmed to invalidate cache blocks in the file system cache including sectors being written to by the uncached write interface (*see column 8, line 36 through column 9, line 39*).

Response to Arguments

7. Applicant's arguments with respect to claims 1-28, 32-54, and 58-73 have been considered but are moot in view of the new ground(s) of rejection.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to JARED M. BIBBEE whose telephone number is (571)270-1054. The examiner can normally be reached on IFP.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Apu Mofiz can be reached on 571-272-4080. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/J. M. B./
Examiner, Art Unit 2161

/Apu M Mofiz/
Supervisory Patent Examiner, Art Unit 2161